Utilizing Green Energy Prediction to Schedule Mixed Batch and Service Jobs in Data Centers

Baris Aksanli (Presenter)
Jagannathan Venkatesh
Liuyi Zhang
Tajana Rosing
Computer Science and Engineering Department
University of California, San Diego

INTERNATIONAL WORKSHOP ON ENVIRONMENT AND ENERGY 2011

11/16/2011

This work was sponsored in part by the Multiscale Systems Center (MuSyC), National Science Foundation (NSF) Project GreenLight, NSF ERC CIAN, NSF Variability Expedition, NSF Flash Gordon, CNS, NSF IRNC TransLight/StarLight, Oracle, Microsoft and Google.

Data Center Energy Consumption and Carbon Emission

- Power consumption of data centers increase rapidly
  - Millions of MWh, reflected as billions of dollars in the electricity bill
- Research and investment on reducing the power consumption
  - Efficient resource utilization
  - Better cooling strategies
  - Better power infrastructure
- Power efficiency increases, but total consumption does not decrease
  - A lot of new data centers are being deployed
  - Recent study shows high power consumption leads to significant carbon emission
  - If no action is taken, data center carbon emission will further increase

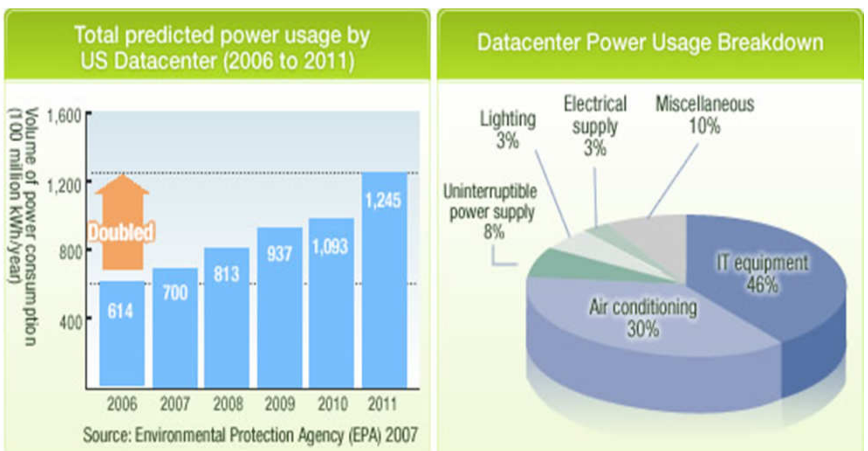

Data center power consumption [Hitachi Eco-friendly Data center project, http://www.hitachi.com]

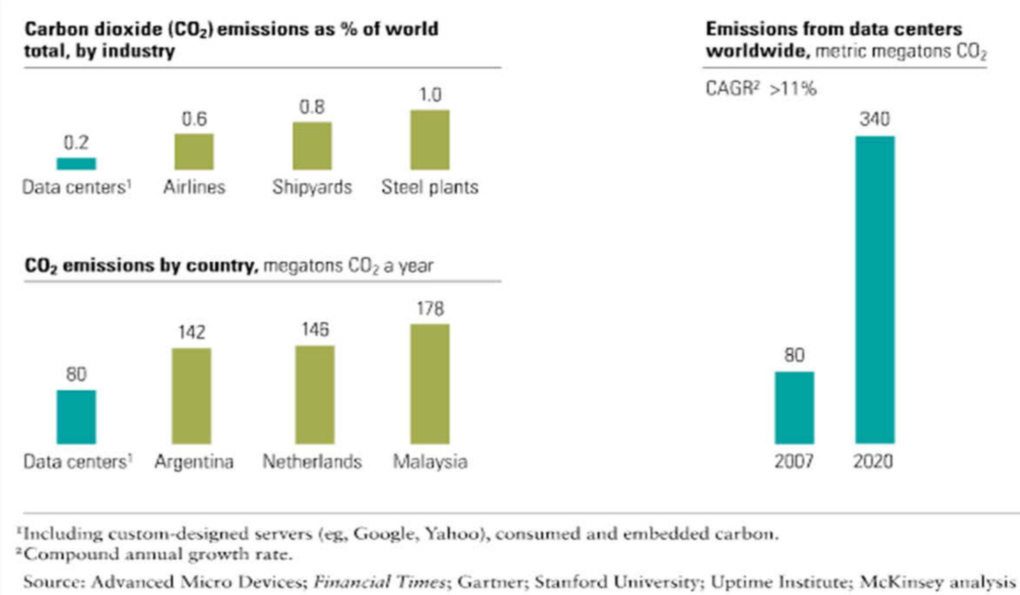

Key issues for distributed renewable-powered datacenters

- Green energy availability varies dramatically
  - Instantaneous use may lead to significant energy efficiency losses
  - Prediction is needed
- Datacenter computing requires consistent performance
  - Infrastructure that monitors and manages computation in datacenters has to be aware of performance costs
    - Service response times are around 100ms, Max 10% batch job throughput hit
- Energy costs of datacenters are typically higher than green energy availability
  - Brown energy needs to be present to both supplement green and as "insurance" to meet performance constraints
  - Improvements in computation & networking infrastructure energy efficiency are necessary

Our Contribution

- Mix services and batch jobs while meeting SLAs for both
  - Increases the number of "available" servers to run more jobs
- Green energy used to run additional batch jobs
  - More jobs completed at lower brown energy cost
  - Brown energy supplements green and ensures performance constraints are met regardless of green energy availability
- Our short term prediction algorithms increase the utilization efficiency of renewable energy
  - Use prediction window of 30 min as most MapReduce jobs complete in less than 30 min [CMU09]

[CMU09] S. Kavulya, J. Tan, R. Gandhi and P. Narasimhan. An Analysis of Traces from a Production MapReduce Cluster. Carnegie Mellon University, Technical Report 2009

<- Wind Turbines

Solar Panels ->

Solar Energy Prediction Algorithms

- Algorithms compared:
  - Exponentially Weighted Moving Average (EWMA)
    - $E_{pred}(t) = \alpha * E_{pred}(t-1) + (1-\alpha) * E_{obs}(t-1)$
    - Predicts based on old prediction and old observation
      - α is the weighting factor
  - Extended EWMA (eEWMA)
    - $E_{pred}(t+1) = \beta * E_{obs}(t) * (1+\varepsilon_1) + (1-\beta) * E_{obs}(t-1) * (1+\varepsilon_2)$

$$\varepsilon_1 = \frac{E_{obs}(t) - E_{pred}(t)}{E_{obs}(t)} \qquad \varepsilon_2 = \frac{E_{obs}(t-1) - E_{pred}(t-1)}{E_{obs}(t-1)}$$

    - Predicts based on 2 previous steps
      - β is the weighting factor

For the algorithms above E_{obs} denotes the observation whereas E_{pred} denotes the predicted value

Solar Energy Prediction Algorithms

- Weather Conditioned Moving Average (WCMA)
  - D days and N slots/day
  - Quicker calculation of a baseline and adaptation to seasonal changes
  - The main algorithm:

$$E(d, n+1) = \alpha * E(d, n) + GAP_k * (1 - a) * M_D(d, n+1)$$

where:

$$M_D(d, n) = \frac{\sum_{i=d-1}^{d-D} E(i, n)}{D}$$

α is the weighting factor; d and n represent the day and slot respectively

$E(d, n)$ is the predicted energy for day d and slot n

  - $M_D(d, n)$ the mean of previous D days' values for the slot to be predicted
  - GAP_k weights the mean values by the distance to the point to be predicted

Algorithms tested with 7 days of trace of a large-scale solar panel at UCSD in variable weather conditions

Solar Energy Prediction Results

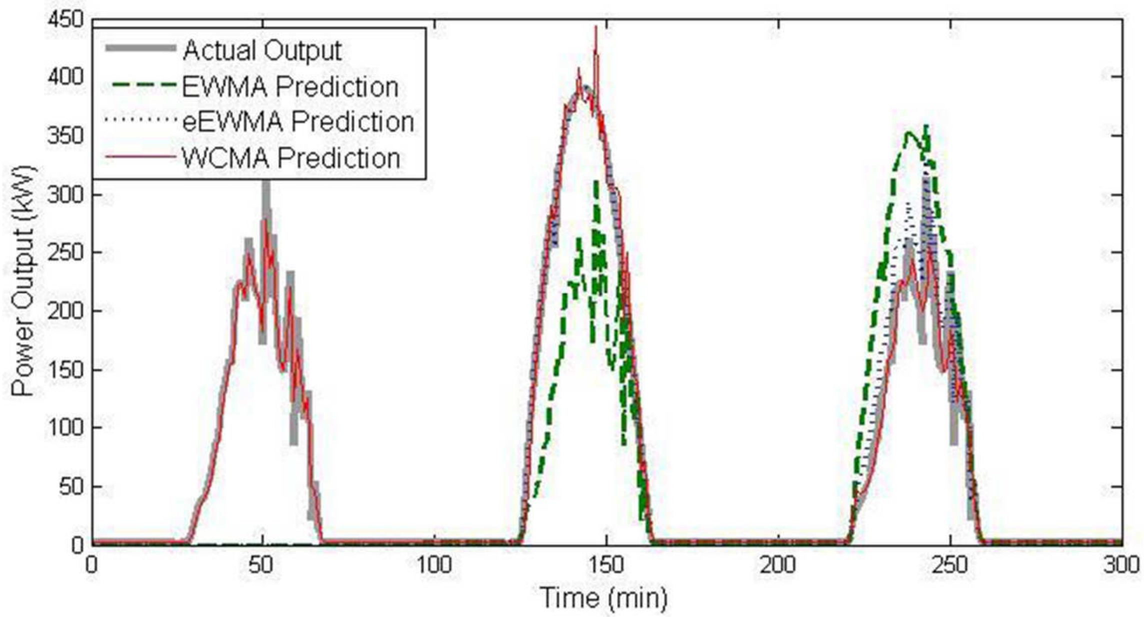

- Prediction values are relatively accurate when they are under consistent conditions
- EWMA -> 32.6 % error
- Extended EWMA -> 23.4% error
- WCMA -> 9.6% error

Wind Energy Prediction Algorithms

- The algorithm we use:
  - Wind velocity-> Wind Power
    - $P_w = 0.5\rho v^3$
      - P is the power, ρ is the air density and v is the wind velocity
  - Combination of a weighted nearest-neighbor (NN) tables and wind power curve models
    - Flexible power curve model, allowing seasonal changes
  - Table update rule:
    - $P_{new}(v, d) = \alpha * P_{obs}(v, d, t) + (1 - \alpha) * P_{old}(v, d)$
  - Prediction:
    - $P_{pred}(v, d, t + k) = P(v(t + k), d(t + k))$

    where v and d are the wind velocity and direction respectively

A. Kusaik, H. Zheng, Z. Song. "Short term prediction of wind farm power: A Data Mining approach". IEEE TEC March 2009.

Wind Energy Prediction Algorithms

- State-of-the-art
  - Time-series analysis algorithms
  - Long term prediction, models have up to 120 parameters
    - Prediction for long periods, e.g. hours, days
- Comparison:

Algorithm	Mean error (%)
SVMReg	19.8
MLP	26.6
M5P Tree	20.6
REP Tree	24.9
Bagging Tree	20.2
NN Predictor	21.2

- NN predictor has good accuracy compared to state-of-the-art TSA algorithms.
  - We use only 2 parameters: speed and direction
  - Same level of accuracy w/ less complexity

Wind Energy Prediction Results

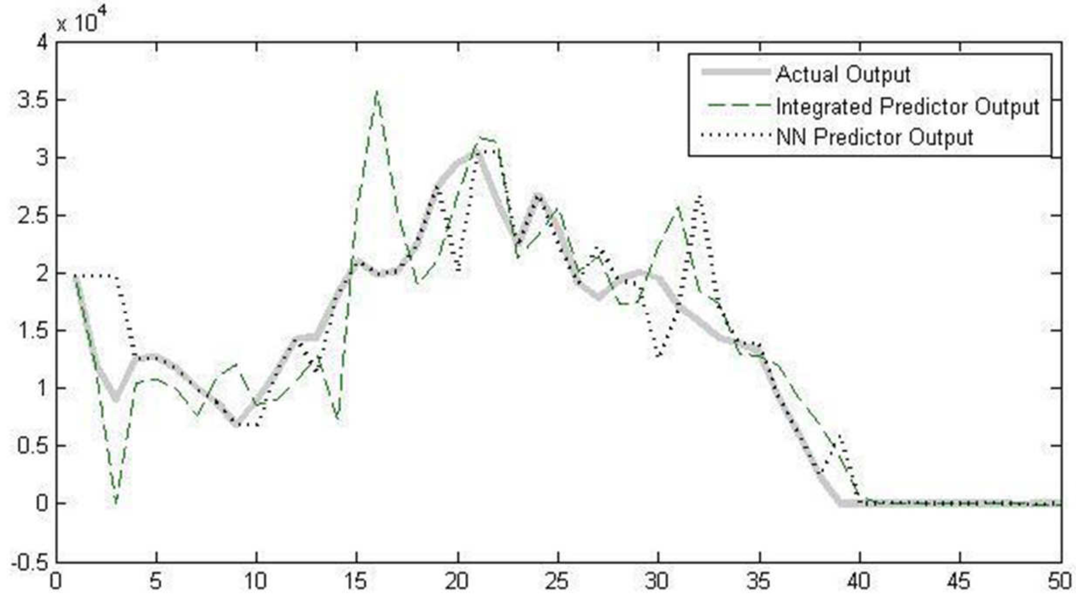

- 21.2% error vs. state-of-the-art integrated predictor 48.2% error

Tested against a wind farm installation over a year's worth of power output data provided by the NREL, and the meteorological data provided by the NCDC

System Model

- Uses estimated green energy to schedule additional batch jobs without reducing the service job response time
  - Initiate more MapReduce subtasks within VMs with available green energy
    - More batch jobs are done but at a slower job completion time
    - End a subtask if the green energy supply level decreases below predicted value
  - Requires understanding of per VM energy costs and the ability to control SLAs at the VM level

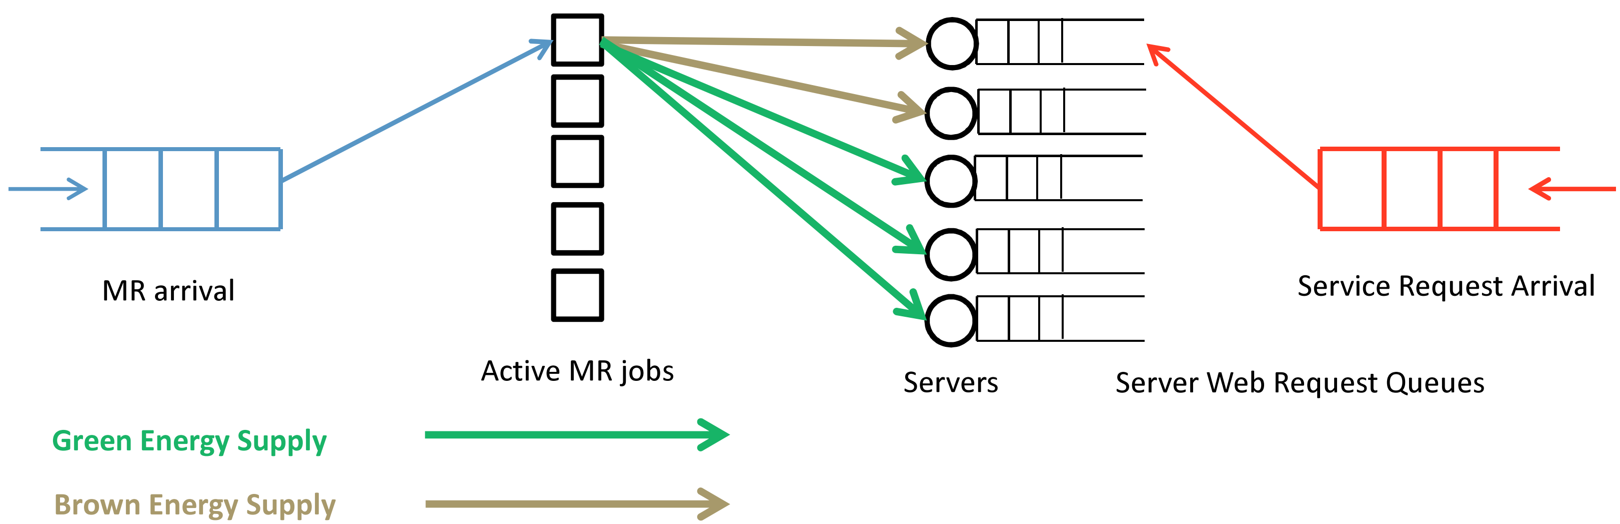

Energy management with virtualization

vGreen

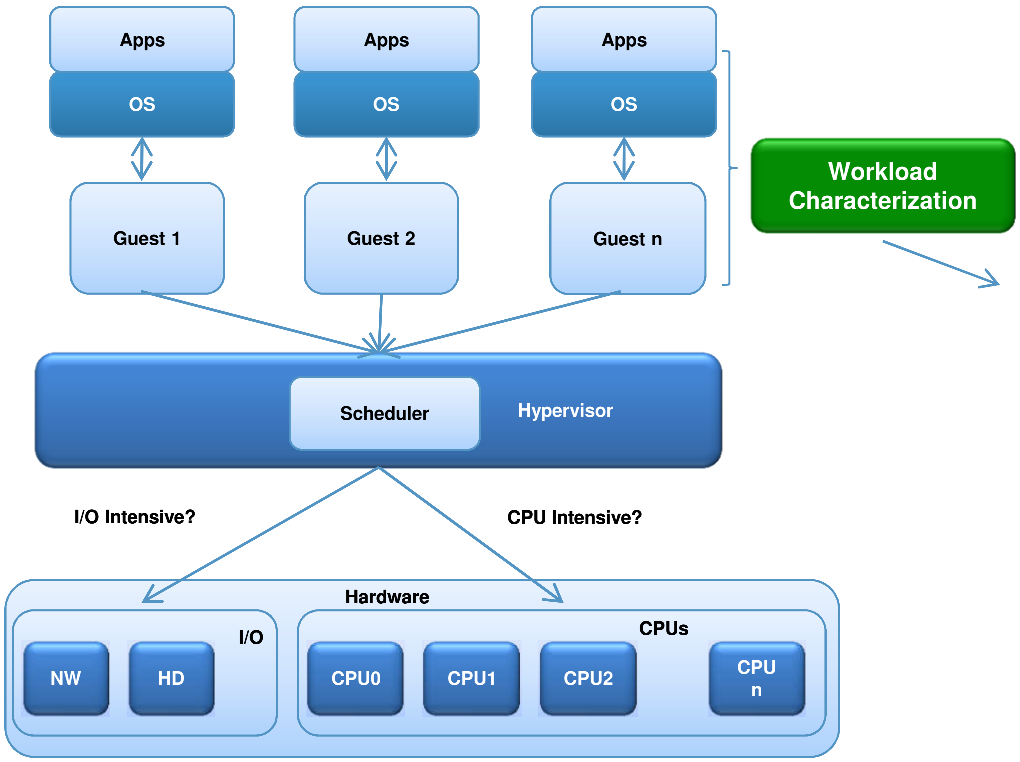

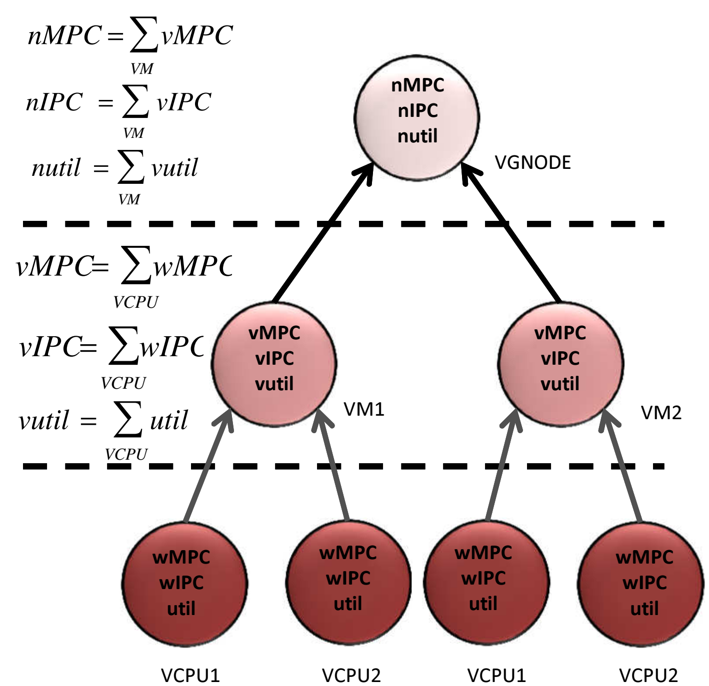

- Scheduling
  - Co-locate guests with orthogonal characteristics
- Management policies
  - Based on the metrics maintained per guest

✓ **Avg 35% Energy Savings**

✓ **Avg 40% speedup**

Gaurav Dhiman, Giacomo Marchetti, and Tajana Rosing. 2009. vGreen: a system for energy efficient computing in virtualized environments. In Proceedings of the 14th ACM/IEEE international symposium on Low power electronics and design (ISLPED '09).

Experimental setup & validation

- Globally distributed datacenters connected
  - 5 datacenters, 12 routers, modeled after ESnet
  - Solar traces from UCSD
  - Wind traces from NREL
- Simulator estimation error <10% measurements

Simulator validation	Measured Value	Simulated Value	Ave. Error
Avg. Power Consumption	246 W	251 W	3%
Rubis QoS ratio	0.08	0.085	6%
Avg. MapReduce Completion Time	112 min	121 min	8%

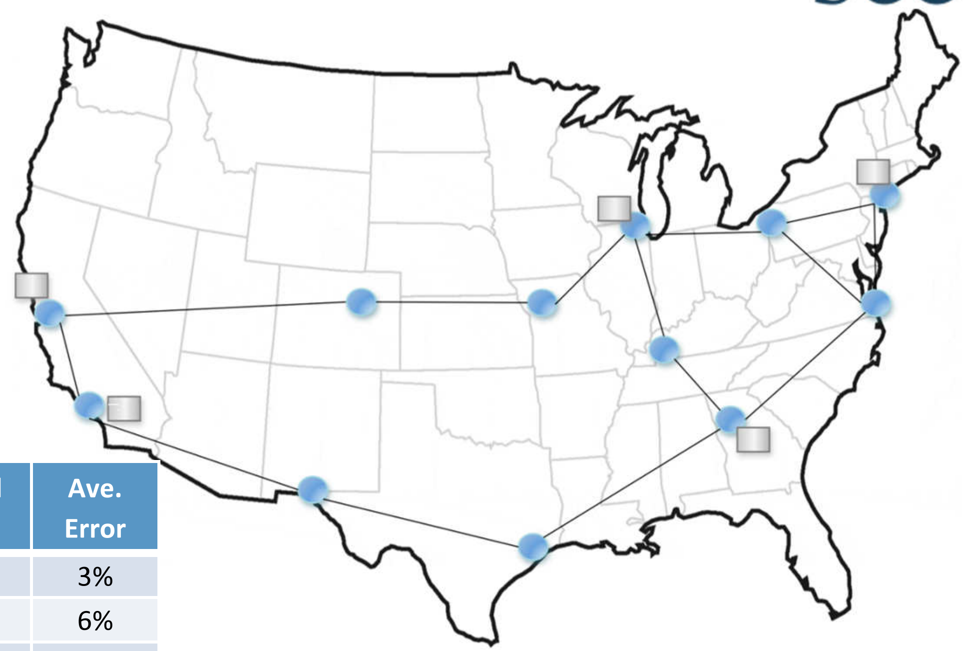

- Jobs run in vGreen VMs on Nehalem servers in datacenter container
  - Rubis used for services with 100ms 90th%ile response time constraint
  - MapReduce used for batch jobs with 10% max job completion time reduction (max 5 cores on Nehalem server)
  - VM migration enabled

Benefits of Green Energy Prediction

- Compare our green energy *predictive* scheduler with *instantaneous* usage of green energy

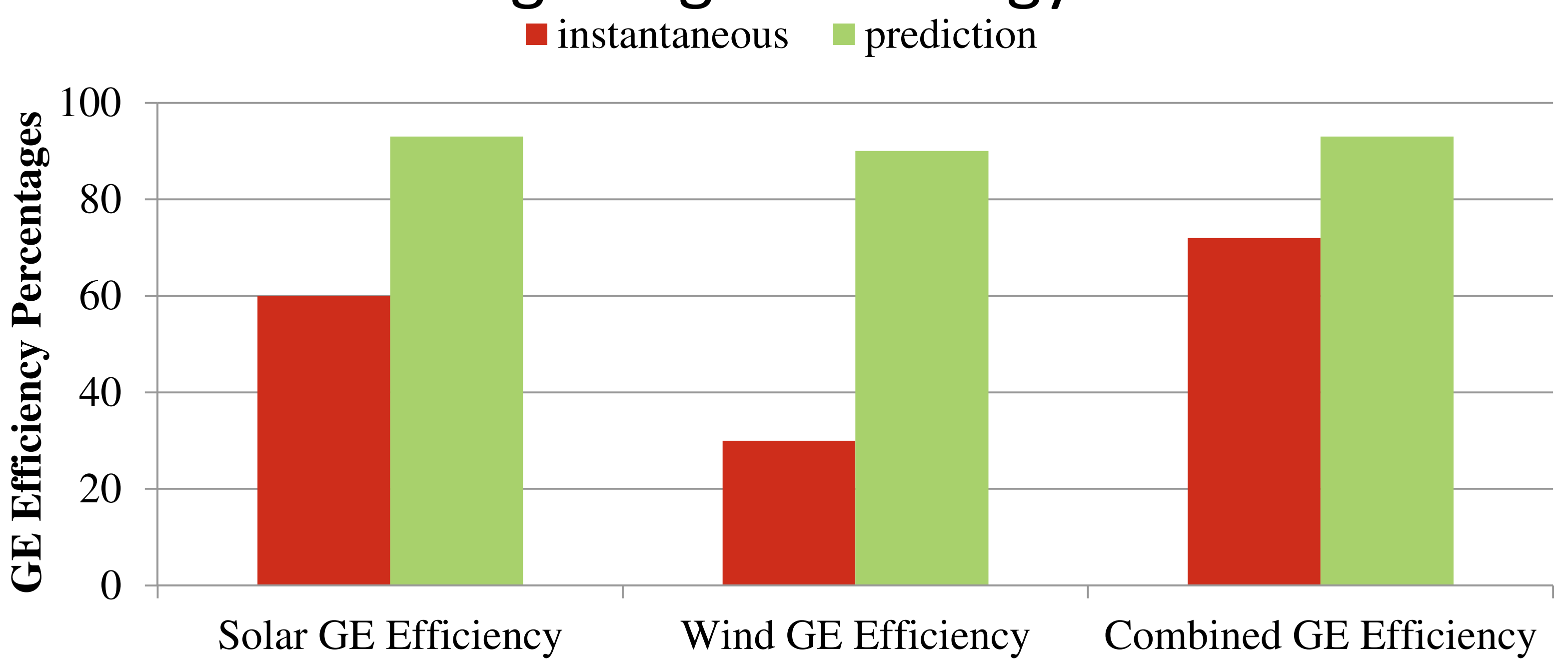

Prediction has 93% green energy efficiency

Green energy efficiency is a ratio of green energy consumed for useful work compared to the total green energy available

Benefits of Green Energy Prediction

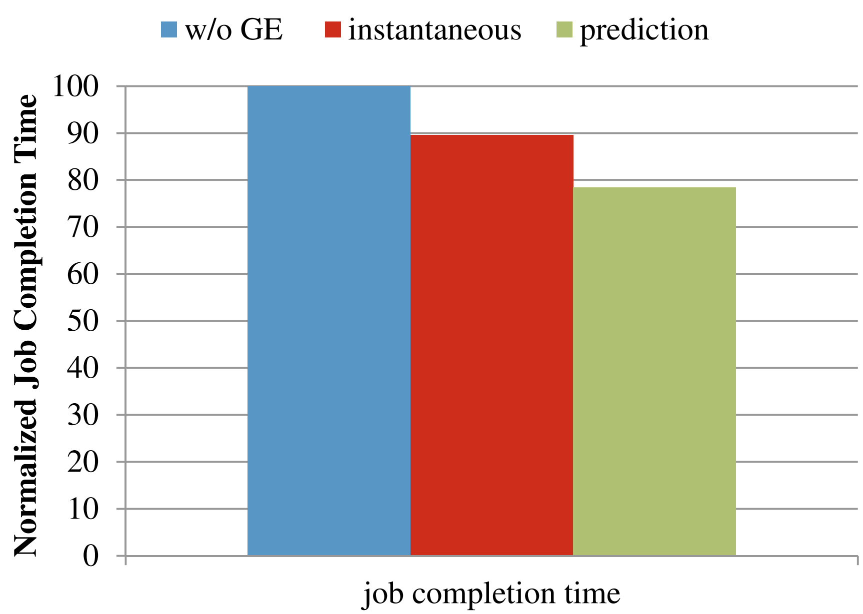

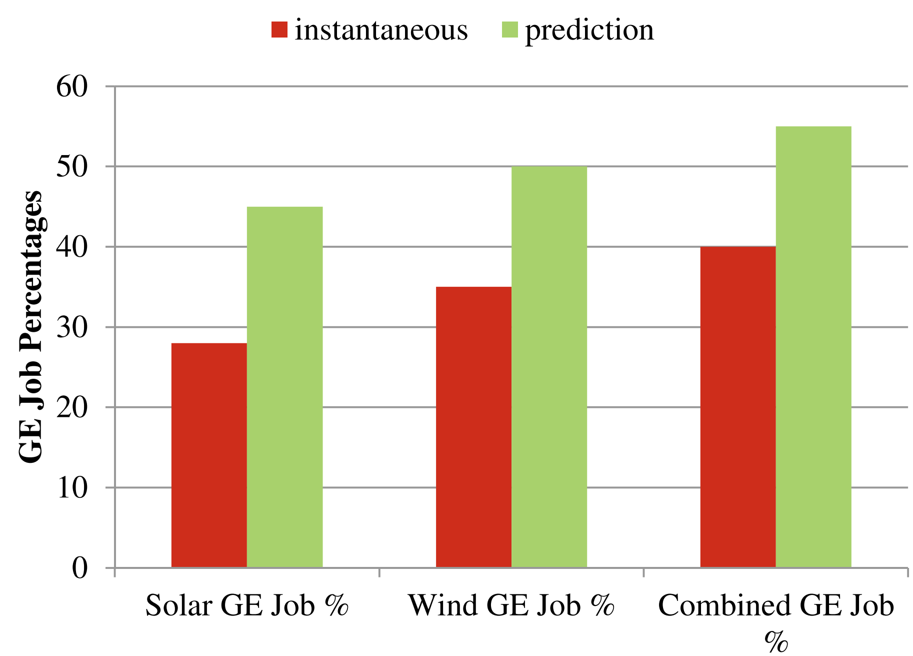

Prediction has 15% faster batch job completion time vs. instantaneous

On average, **5x fewer batch tasks need to be terminated** when using green energy prediction vs. instantaneous usage

Prediction has **~2x more jobs completed** with green energy vs. instantaneous usage

Green energy job %: ratio of batch jobs completed with GE over all jobs.

Summary

- Energy efficiency in data centers is a huge concern
  - Renewable energy can be used
    - To reduce the brown energy costs
    - To reduce the environmental effects by decreasing carbon emissions
- The negative effect of green energy variability can be eliminated by:
  - *Predicting* the green energy
  - Using green energy as the *supplementary* energy source
    - Scheduling additional work with green energy
- Prediction outperforms instantaneous usage:
  - Better performance, 15% lower job completion time
  - 2x more efficient green energy usage
  - 5x fewer task termination

Thanks for your attention!

Questions?

see

System Energy Efficiency Lab
seelab.ucsd.edu

BACKUP SLIDES

Related Work

- Using renewables to reduce "brown" energy consumption in data centers
  - [CNSM10] Using wind and solar energy to cap the power usage of a data center environment
    - Models only service requests
    - The peak brown power is fixed and green energy supplements the rest
      - Stored energy is necessary not to suffer from green energy variability
    - Less stored energy is necessary with additional green energy
      - Both stored and grid energy reduces by 20% with 100 kW PV (200 kW peak)
  - [ICGC10] Optimizing the amount of green energy to be bought based on price
    - Brown and green energy have different market prices
      - Green energy is more expensive, but extra carbon tax for per brown energy
    - Both market prices and carbon tax change over time
    - Minimization over total cost (energy costs + carbon tax)

[CNSM10] D. Gmach, J. Rolia, C. Bash, Y. Chen, T. Christian, A. Shah, R. Sharma and Z. Wang. "Capacity Planning and Power Management to Exploit Sustainable Energy"
[ICGC10] K. Le, R. Bianchini, T. D. Nguyen, O. Bilgir, M. Martonosi. "Capping the brown energy consumption of Internet services at low cost"

Related Work

- [ISSST10] Augmenting a data center with PV and municipal solid waste based energy
  - The latter is dominant and designed as static
    - 15 kW solid waste energy + 4 kW PV is enough for a data center of ~200 servers
  - Does not model variability of green energy
    - The dominant energy type is static and it is ≥ 80% of overall energy
- [IEEE11] Using green energy to execute MapReduce jobs
  - Aligning job execution with green energy availability
  - Using instantaneous green energy
    - Models variability, but does not propose a solution for its adverse effects

[ISSST10] D. Gmach, Y. Chen, A. Shah, J. Rolia, C. Bash, T. Christian, R. Sharma. "Profiling sustainability of data centers"
[IEEE11] A. Krioukov, C. Goebel, S. Alspaugh, Y. Chen, D. Culler, R. Katz. "Integrating Renewable Energy Using Data Analytics Systems: Challenges and Opportunities." *IEEE Data Engineering Bulletin. March 2011.*

Simulation Setup

- Simulated a data center with 4 days of solar and wind traces
  - Both prediction and instantaneous values
- Parameters used:

Parameter	Value	Parameter	Value
Mean Web Request Inter-arrival time	5ms	**Mean Web Request Service time**	12ms
Service Request SLA	150 ms	**Mean MapReduce Task Service time**	10 min
Mean MapReduce Job Inter-arrival time	2 min	**Idle Server Power**	212.5W
Peak Server Power	312.5W	**Number of servers**	200